

=====

Paper ID: #359

=====

Paper Title: Tractable Pareto Optimization of Temporal Preferences

Subtitle:

Keyword 1: constraint satisfaction

Keyword 2: temporal reasoning

Keyword 3:

ABSTRACT (200 words): This paper focuses on temporal constraint problems where the objective is to optimize a set of local preferences for when events occur. In previous work, a subclass of these problems has been formalized as a generalization of Temporal CSPs, and a tractable strategy for optimization has been proposed, where global optimality is defined as maximizing the minimum of the component preference values. This criterion for optimality, which we call "Weakest Link Optimization" (WLO), is known to have limited practical usefulness because solutions are compared only on the basis of their worst value; thus, there is no requirement to improve the other values. To address this limitation, we introduce a new algorithm that re-applies WLO iteratively in a way that leads to improvement of all the values. We show the value of this strategy by proving that, with suitable preference functions, the resulting solutions are Pareto Optimal.

Contact Author: Robert Morris

Affiliation 1: NASA Ames Research Center

Affiliation 2:

Street Address1: 269-3

Street Address2:

City/Town: Moffett Field

State: California

Zip: 94035

Country: US

Phone:

Fax:

Email: morris@email.arc.nasa.gov

Co-Author 1: Paul Morris pmorris@email.arc.nasa.gov

Co-Author 2: Lina Khatib lina@email.arc.nasa.gov

Co-Author 3: Brent Venable kvenable@math.unipd.it

Co-Author 4:

Co-Author 5:

Co-Author 6:

Co-Author 7:

Co-Author 8:

Co-Author 9:

Co-Author 10:

The author(s) certifies that this paper has not already been accepted by nor is it currently under review for a journal or another conference, nor will it be submitted for such during IJCAI-03's review period.

Tractable Pareto Optimization of Temporal Preferences

Content Areas : Constraint Satisfaction, Temporal Reasoning

Abstract

This paper focuses on temporal constraint problems where the objective is to optimize a set of local preferences for when events occur. In previous work, a subclass of these problems has been formalized as a generalization of Temporal CSPs, and a tractable strategy for optimization has been proposed, where global optimality is defined as maximizing the minimum of the component preference values. This criterion for optimality, which we call “Weakest Link Optimization” (WLO), is known to have limited practical usefulness because solutions are compared only on the basis of their worst value; thus, there is no requirement to improve the other values. To address this limitation, we introduce a new algorithm that re-applies WLO iteratively in a way that leads to improvement of all the values. We show the value of this strategy by proving that, with suitable preference functions, the resulting solutions are Pareto Optimal.

1 Introduction

The notion of *softness* has been used to describe either a constraint or planning goal, indicating that either can be satisfied to matters of degree. For example, in an earth orbiting spacecraft, sensitive instruments like imagers have *duty cycles*, which impose restrictions on the amount of use of the instrument. A duty cycle is typically a complex function based on both the expected lifetime of the instrument, as well as short term concerns such as the amount of heat it can be exposed to while turned on. Duty cycles impose constraints on the duration of the periods for which the instrument can be on, but it is natural to view this duration as flexible. For example, this restriction might be waived to capture an important event such as an active volcano. Thus, the flexibility of the duty cycle “softens” the constraint that the instrument cannot be on beyond a certain duration. Reasoning about soft constraints for planning or scheduling is for the purpose of finding a solution that satisfies the constraints to the highest degree possible.

For temporal reasoning problems, a simple method for evaluating the global temporal preference of a solution to a Temporal CSP involving local temporal preferences was

introduced in [Khatib *et al.*, 2001], based on maximizing the minimally preferred local preference for a time value. Because the locally minimally preferred assignment can be viewed as a sort of “weakest link” with respect to the global solution, we dub this method “weakest link optimization” (WLO), in the spirit of the television game show. WLO was chosen primarily for reasons of computational efficiency. Specifically, WLO can be formalized using a generalization of Simple Temporal Problems (STPs), called STPs with Preferences (STPPs), that preserves the capability to tractably solve for solutions (with suitable preference functions associated with the temporal constraints). Unfortunately, as often occurs, this efficiency has a price. Specifically, WLO offers an insufficiently fine-grained method for comparing solutions, for it is based on a single value, viz., the “weakest link.” It is consequently easy to conceive of examples where WLO would accept intuitively inferior solutions because of this myopic focus. Although it is possible to consider more robust alternatives to a WLO strategy for evaluating solutions, it is not clear whether any of these methods would preserve the computational benefits of WLO. This impasse is the starting point of the work described in this paper.

We propose here an approach to making WLO more robust by combining it with an iterative strategy for solving STPPs. The process involves repeatedly restricting temporal values for the weakest links, resetting their preference values, and applying the WLO procedure to the reduced problem that results from these changes. The intuition is a simple one, and we motivate this technique with an example from a Mars Rover planning domain. In Section 2, we summarize the soft constraint problem solver based on WLO introduced previously. We then illustrate in section 3 the deficiencies of WLO on a simple example, which also reveals the intuition underlying the proposed strategy for overcoming this deficiency. The main contribution of this paper is discussed in sections 4 and 5, which formalize this strategy and prove that any solution generated by an application of this strategy is in the set of Pareto optimal solutions for the original problem.

2 Reasoning about preferences with soft constraints

This section reviews the material first presented in [Khatib *et al.*, 2001]. There, a class of constrained optimization prob-

lems, called Temporal Constraint Satisfaction Problems with Preferences (TCSPPs), was first defined. A TCSPP is a generalization of classical TCSPs which allows for a representation of soft constraints. In classical TCSPs [Dechter *et al.*, 1991], a unary constraint over a variable X representing an event restricts the domain of X , representing its possible times of occurrence; the constraint is then shorthand for $(a_1 \leq X \leq b_1) \vee \dots \vee (a_n \leq X \leq b_n)$. Similarly, a binary constraint over X and Y restricts the values of the distance $Y - X$, in which case the constraint can be expressed as $(a_1 \leq Y - X \leq b_1) \vee \dots \vee (a_n \leq Y - X \leq b_n)$. A uniform, binary representation of all the constraints results from introducing a variable X_0 for the *beginning of time*, and recasting unary constraints as binary constraints involving the distance $X - X_0$.

A *soft temporal constraint* is a pair $\langle I, f \rangle$, where I is a set of intervals $\{[a, b], a \leq b\}$ and f is a function from $\bigcup I$ to a set A of preference values. To compare and combine values from this set, A is organized in the form of a *c-semiring* [Bistarelli *et al.*, 1997]. A *semiring* is a tuple $\langle A, +, \times, 0, 1 \rangle$ such that

- A is a set and $0, 1 \in A$;
- $+$, the additive operation, is commutative, associative and 0 is its identity element ($a + 0 = a$);
- \times , the multiplicative operation, is associative, distributes over $+$, 1 is its identity element and 0 is its absorbing element ($a \times 0 = 0$).

A *c-semiring* is a semiring in which $+$ is idempotent (i.e., $a + a = a, a \in A$), 1 is its absorbing element, and \times is commutative. The class of TCSPPs in which each constraint consists of a single interval is called *Simple Temporal Problems with Preferences* (STPPs).

A *solution* to a TCSPP is a complete assignment to all the variables that satisfies the temporal constraints. An arbitrary assignment of values to variables has a *global preference value*, obtained by combining the local preference values using the semiring operations. A *c-semiring* induces a partial order relation \leq_S over A to compare preference values of arbitrary assignments; $a \leq_S b$ can be read *b is more preferred than a*. Classical Temporal CSPs can be seen as a special case of TCSPP, with “soft” constraints that assign the “best” preference value to each element in the domain, and the “worst” value to everything else. The optimal solutions of a TCSPP are those solutions which have the best preference value in terms of the ordering \leq_S .

Weakest Link Optimization (WLO) is formalized via the semiring $S_{WLO} = \langle A, \max, \min, 0, 1 \rangle$, that is, where for $a, b \in A$, $a + b = \max(a, b)$ and $a \times b = \min(a, b)$, and 1 (0) is the best (worst) preference value. Given a solution t in a TCSPP with semiring S_{WLO} , let $T_{ij} = \langle I_{ij}, f_{ij} \rangle$ be a soft constraint over variables X_i, X_j and (v_i, v_j) be the projection of t over the values assigned to variables X_i and X_j (abbreviated as $(v_i, v_j) = t_{\downarrow X_i, X_j}$). The corresponding preference value given by f_{ij} is $f_{ij}(v_j - v_i)$, where $v_j - v_i \in I_{ij}$. The global preference value of t , $val(t)$, is defined as $val(t) = \min\{f_{ij}(v_j - v_i) \mid (v_i, v_j) = t_{\downarrow X_i, X_j}\}$. Thus, a “weakest link value” for a solution t is any minimum




Figure 1: “Chopping” a semi-convex function

$f_{ij}(v_j - v_i)$ that determines $val(t)$, and the *WLO-optimal solutions* to a problem are the ones that have a maximum weakest link value.

As with classical (binary) CSPs, TCSPPs can be arranged to form a network of nodes representing variables, and edges labeled with constraint information. Given a network of soft constraints, under certain restrictions on the properties of the semiring, it can be shown that local consistency techniques can be applied in polynomial time to find an equivalent minimal network in which the constraints are as explicit as possible. The restrictions that suffice for this result apply to

1. the “shape” of the preference functions used in the soft constraints;
2. the multiplicative operator \times (it should be idempotent); and
3. the ordering of the preference values (\leq_S must be a total ordering).

The class of restricted preference functions that suffice to guarantee that local consistency can be meaningfully applied to soft constraint networks is called *semi-convex*. This class includes linear, convex, and also some step functions. All of these functions have the property that if one draws a horizontal line anywhere in the Cartesian plane of the graph of the function, the set of X such that $f(X)$ is not below the line forms an interval. Semi-convexity is preserved under the operations performed by local consistency (intersection and composition). STPPs with semiring S_{WLO} can easily be seen to satisfy these restrictions.

The same restrictions that allow local consistency to be applied are sufficient to prove that STPPs can be solved tractably. Finding an optimal solution of the given STPP with semi-convex preference functions reduces to a two-step search process consisting of iteratively choosing a preference value, “chopping” every preference function at that point, then solving a STP defined by considering the interval of temporal values whose preference values lies above the chop line (semi-convexity ensures that there is a single interval above the chop point, hence that the problem is indeed an STP). Figure 1 illustrates the chopping process. It has been shown that the “highest” chop point that results in a solvable STP in fact produces an STP whose solutions are exactly the optimal solutions of the original STPP. Binary search can be used to select candidate chop points, making the technique for solv-

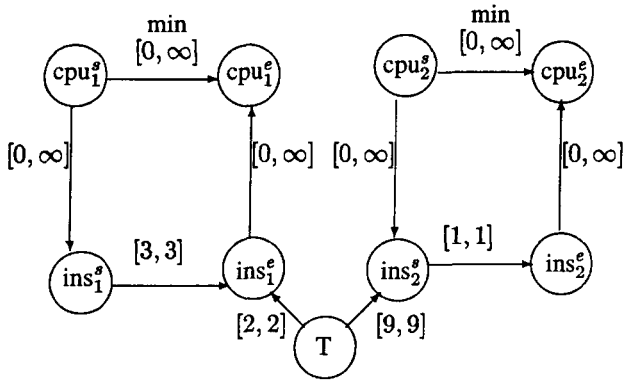


Figure 2: The STPP for the Rover Science Planning Problem where T is any timepoint

ing the STPP tractable. The second step, solving the induced STP, can be performed by transforming the graph associated with this STP into a distance graph, then solving two single-source shortest path problems on the distance graph. (The solutions to these provide upper and lower time bounds for each event. If the problem has a solution, then for each event it is possible to arbitrarily pick a time within its time bounds, and find corresponding times for the other events such that the set of times for all the events satisfy the interval constraints.) For N nodes and E edges, the complexity of this phase is $O(EN)$ (using the Bellman-Ford algorithm [Cormen *et al.*, 1990]).

3 The problem with WLO

Formalized in this way, WLO offers a coarse method for comparing solutions, one based on the minimal preference value over all the projections of the solutions to local preference functions. Consequently, the advice given to a temporal solver by WLO may be insufficient to find solutions that are intuitively more globally preferable. For example, consider the following simple Mars rover planning problem, illustrated in Figure 2. The rover has a sensing instrument and a CPU. There are two sensing events, of durations 3 time units and 1 time unit (indicated in the figure by the pairs of nodes labeled ins_1^s, ins_1^e and ins_2^s, ins_2^e respectively). There is a hard temporal constraint that the CPU be on while the instrument is on, as well as a soft constraint that the CPU should be on as little as possible, to conserve power. This constraint is expressed in the STPP as a function from temporal values indicating the duration that the CPU is on, to preference values. For simplicity, we assume that the preference function on the CPU duration constraints is the negated identity function, $f(t) = -t$; thus higher preference values, i.e. shorter durations, are preferred. Because the CPU must be on at least as long as the sensing events, any globally preferred solution using WLO has preference value -3. The set of solutions that have the optimal value includes solutions in which the CPU duration for the second sensing event varies from 1 to 3 time units. The fact that WLO is unable to discriminate between the global values of these solutions, despite the fact that the one with 1 time unit is obviously preferable to the others, is a

clear limitation of WLO.

One way of formalizing this drawback of WLO is to observe that a WLO policy is not *Pareto Optimal*. To see this, we reformulate the set of preference functions of a STPP, f_1, \dots, f_m as criteria requiring simultaneous optimization, and let $s = [t_1, \dots, t_n]$ and $s' = [t'_1, \dots, t'_m]$ be two solutions to a given STPP. s' dominates s if for each j , $f_j(t_j) \leq f_j(t'_j)$ and for some k , $f_k(t_k) < f_k(t'_k)$. In a Pareto optimization problem, the *Pareto optimal* set of solutions is the set of non-dominated solutions. Similarly, let the *WLO-optimal* set be the set of optimal solutions that result from applying the chopping technique for solving STPPs described above. Clearly, applying WLO to an STPP does not guarantee that the set of WLO-optimal solutions is a Pareto optimal set. In the rover planning problem, for example, suppose we consider only solutions where the CPU duration for the first sensing event is 3. Then the solution in which the CPU duration for the second sensing event is 1 time unit dominates the solution in which it is 2 time units, but both are WLO-optimal, since they have the same weakest link value.¹

Assuming that Pareto-optimality is a desirable objective in optimization, a reasonable response to this deficiency is to replace WLO with an alternative strategy for evaluating solution tuples. A natural, and more robust alternative evaluates solutions by summing the preference values, and ordering them based on preferences towards larger values. (This strategy would also ensure Pareto optimality, since every maximum sum solution is Pareto optimal.) This policy might be dubbed "utilitarian." The main drawback to this alternative is that the ability to solve STPPs tractably is no longer apparent. The reason is that the formalization of utilitarianism as a semiring forces the multiplicative operator (in this case, *sum*), not to be idempotent (i.e., $a + a \neq a$), a condition required in the proof that a local consistency approach is applicable to the soft constraint reasoning problem.

Of course, it is still possible to apply a utilitarian framework for optimizing preferences, using either local search or a complete search strategy such as branch and bound. Rather than pursuing this direction of resolving the problems with WLO, we select another approach, based on an algorithm that interleaves flexible assignment with propagation using WLO.

4 An algorithm for Pareto Optimization

The proposed solution is based on the intuition that if a constraint solver using WLO could iteratively "ignore" the weakest link values (i.e. the values that contributed to the global solution evaluation) then it could eventually recognize solutions that dominate others in the Pareto sense. For example, in the Rover Planning problem illustrated earlier, if the weakest link value of the global solution could be "ignored," the WLO solver could recognize that a solution with the CPU on for 1 time unit during the second instrument event is to be preferred to one where the CPU is on for 2 or 3 time units.

We formalize this intuition by a procedure wherein the original STPP is transformed by iteratively selecting what we

¹This phenomenon is often referred to in the literature as the "drowning effect."

Inputs: an STPP $P = (V, C)$

Output:

A STP (V, C_P) whose solutions are Pareto optimal for P .

(1) $C_P = C$

(2) Do

(3) Solve (V, C_P) using WLO

(4) Delete all weakest link soft constraints from C_P

(5) For each deleted constraint $\langle [a, b], f \rangle$,

(6) add $\langle [a_{opt}, b_{opt}], f_{best} \rangle$ to C_P

(7) until (3)-(6) leave (V, C_P) unchanged

(8) Return (V, C_P)

Figure 3: STPP solver WLO+ returns a solution in the Pareto optimal set of solutions

shall refer to as a *weakest link constraint*, changing the constraint in such a way that it can effectively be “ignored,” and solving the transformed problem. A weakest link (soft) constraint for a WLO set of solutions is one in which the preference value of its duration in all the WLO solutions is the same as the chop level v of the optimal STP using WLO. For example, after applying WLO to the problem in Figure 2, the CPU duration constraint associated with the first sensing event will be a weakest link, since it now has a fixed preference value of -3. However, the CPU constraint for the second event will not be a weakest link since its preference value can still vary from -3 to -1.

We also define a weakest link constraint to be *open* if v is not the “best” preference value (i.e., $v < 1$, where 1 is the designated “best” value among the values in the semi-ring).

Formalizing the process of “ignoring” weakest link values is a two-step process of restricting the weakest links to their intervals of optimal temporal values, while eliminating their WLO restraining influence by resetting their preferences to a single, “best” value. Formally, the process consists of:

- Squeezing the temporal domain to include all and only those values which are optimally preferred; and
- Replacing the preference function by one that assigns the most preferred value (i.e. 1) to each element in the new domain.

The first step ensures that only the best temporal values are part of any solution, and the second step allows WLO to be re-applied to eliminate Pareto-dominated solutions from the remaining solution space.

The algorithm WLO+ (Figure 3) returns a Simple Temporal Problem (STP) whose solutions are contained in the WLO-optimal, Pareto-optimal solutions to the original STPP, P . Where C is a set of soft constraints, the STPP (V, C_P) is solved (step 3) using the chopping approach described earlier. In step 5, we denote the soft constraint that results from the two-step process described above as $\langle [a_{opt}, b_{opt}], f_{best} \rangle$, where $[a_{opt}, b_{opt}]$ is the interval of temporal values that are optimally preferred, and f_{best} is the preference function such that $f_{best}(v) = 1$ for any input value v . Notice that the run time of WLO+ is $O(|C|)$ times the time it takes to execute $Solve(V, C_P)$, which is a polynomial.




Figure 4: Relationships between Solution Spaces for STPPs that are WLO or Pareto Optimal

We now proceed to prove the main result, in two steps. In this section we assume the existence of weakest links at every iteration of the WLO+ algorithm, and show that the subset of solutions of the input STPP returned by WLO+ is contained in the intersection of WLO-optimal and Pareto-optimal solutions. In the next section we show that, given additional restrictions on the shape of the preference functions, such weakest links can be shown to always exist.

Given an STPP P , let $Sol_{WLO}(P)$ (resp. $Sol_{PAR}(P)$) be the set of WLO-optimal (respectively, Pareto-Optimal) solutions of P , and let $Sol_{WLO+}(P)$ be the set of solutions to P returned by WLO+. Then the result can be stated as follows.

Theorem 1 *If a weakest link constraint is found at each stage of WLO+, then $Sol_{WLO+}(P) \subseteq Sol_{WLO}(P) \cap Sol_{PAR}(P)$. Moreover, if P has any solution, then $Sol_{WLO+}(P)$ is nonempty.*

Proof:

First note that after an open weakest link is processed in steps (4) to (6), it will never again be an open weakest link (since its preference is reset to f_{best}). Since the theorem assumes a weakest link constraint is found at each stage of WLO+, the algorithm will terminate when the weakest link constraint is not open, i.e., when all the soft constraints in C_P have WLO preferences that equal the best (1) value.

Now assume $s \in Sol_{WLO+}(P)$. Since the first iteration reduces the set of solutions of (V, C_P) to $Sol_{WLO}(P)$, and each subsequent iteration either leaves the set unchanged or reduces it further, it follows that $s \in Sol_{WLO}(P)$. Now suppose $s \notin Sol_{PAR}(P)$. Then s must be dominated by a Pareto optimal solution s' . Let c be a soft constraint in C for which s' is superior to s . Thus, the preference value of the duration assigned by s to c cannot be 1. It follows that at some point during the course of the algorithm, c must become an open weakest link. Since s is in $Sol_{WLO+}(P)$, it survives until then, and so it must provide a value for c that is equal to the chop level. However, since s' dominates s , s' must also survive until then. But this contradicts the assumption that c is a weakest link constraint, since s' has a value greater than the WLO chop level. Hence, s is in $Sol_{PAR}(P)$, and so in $Sol_{WLO}(P) \cap Sol_{PAR}(P)$.

Next suppose the original STPP P has at least one solution. To see that $Sol_{WLO+}(P)$ is nonempty, observe that the modifications in steps (4) to (6), while stripping out solutions that are not WLO optimal with respect to (V, C_P) , do retain all the WLO optimal solutions. Clearly, if there is any solution, there is a WLO optimal one. Thus, if the (V, C_P) in




Figure 5: A unique WLO+ Solution.

any iteration has a solution, the (V, C_P) in the next iteration will also have a solution. Since we are assuming the first $(V, C_P) (= (V, C))$ has a solution, it follows by induction that $Sol_{WLO+}(P)$ is nonempty. \square

The theorem shows that it is possible to maintain the tractability of WLO-based optimization while overcoming some of the restrictions it imposes. In particular, it is possible to improve the quality of the flexible solutions generated within an STPP framework from being WLO optimal to being Pareto optimal.

Although the algorithm determines a nonempty set of solutions that are both WLO optimal and Pareto optimal, the set might not include all such solutions. Consider the example in figure 5. Assume the preference function for all soft constraints is given by $f(t) = t$, i.e., longer durations are preferred (signified by the *max* label on the edges). The WLO+ algorithm will retain a single solution where BC and CD are both 5. However, the solution where BC = 2 and CD = 8, which is excluded, is also both Pareto optimal and WLO optimal. (Note that AB, with a fixed value of 1, is the weakest link.)

Many optimization schemes seek what is known as *utilitarian* optimality, where the objective is to maximize the sum of the local preferences. However, the WLO+ solutions are not necessarily utilitarian optimal with respect to all solutions or even the WLO solutions. For example, in figure 5, if the preference function is $f(t) = t^2$, a utilitarian optimal WLO solution would be given by BC = 1 and CD = 9, but WLO+ will still return the solution where BC and CD are both 5.

We can summarize the position taken in this paper by saying that utilitarian strategies, while attractive in many ways, are apparently intractable. The WLO+ approach provides some of the same benefit at lower cost. For example, non-competing constraints are fully optimized by WLO+. For competing constraints, WLO+ tends to divide the preferences as equally as possible. In some applications, this might be more desirable than a utilitarian allocation.

5 Existence of Weakest Links

In this section we show that under suitable conditions, a weakest link constraint always exists. This involves a stronger requirement than for WLO: the preference functions must be continuous and convex, not merely semi-convex. This would include linear functions, cycloids, and upward-pointing parabolas, for example, but not Gaussian curves, or step functions.

The existence theorem depends on a technical lemma regarding STPs. Before proceeding, we note that while a solution s of an STP P is defined in terms of an assignment to each variable, it also determines a value $s(e)$ for each edge e ,

given by $s(e) = s(Y) - s(X)$ where X and Y are the start and end variables of e , respectively. We will use this notation in what follows.

Now consider any consistent STP P . The *minimal network* [Dechter *et al.*, 1991] corresponding to P is another STP P' . The constraints between any two points X and Y in P' are formed by intersecting the constraints induced by all possible paths between X and Y in P . Suppose $[a, b]$ is the constraint on an edge e between X and Y in P' . If $a = b$, we say e is a *rigid* edge; otherwise it is a *flexible* edge. Notice that all solutions must agree on the values they associate with rigid edges. However, we wish to show that it is possible to find two solutions that differ on every one of the flexible edges.

Lemma 1 *Let P be any consistent STP with associated minimal network P' . Suppose s is an arbitrary solution to P . Then there is another solution \hat{s} such that $s(e) \neq \hat{s}(e)$ for each flexible edge e . Moreover, \hat{s} can be made as close to s as we please, i.e., for every $\epsilon > 0$, \hat{s} can be chosen so that $|\hat{s}(e) - s(e)| < \epsilon$ for every edge e .*

Proof: Due to lack of space, we can only outline the proof here. For each flexible edge e , it is easy to find a solution s_e such that $s_e(e) \neq s(e)$. Then, by choosing suitable convex linear combinations of the $\{s_e\}$, one can construct a solution that differs from s on all of the flexible edges. A further convex linear combination with s can be used to produce such a solution arbitrarily close to s . \square

In the following, a preference function f is said to be *convex* if $\{< x, y > | y \leq f(x)\}$ is a convex set. The claim of the existence of weakest links can be stated as follows:

Theorem 2 *Suppose P is an STPP with continuous domains and continuous convex preference functions. Then there will be at least one weakest link constraint for the WLO optimal set of solutions.*

Proof: Consider the (minimal) STP P_{opt} that corresponds to the optimal chopping level for P (as described in the WLO algorithm). Note that for every solution t to P_{opt} , $f(t(e)) \geq opt$ for every edge e , where f is the preference function for the edge.² Moreover, there must be at least one edge e where $f(t'(e)) = opt$ for some solution t' ; otherwise the *opt* level could be increased. Call the set of such edges *candidate* weakest links. Note that if any of the candidate weakest links e is a rigid edge, then $f(t(e)) = f(t'(e)) = opt$ for all solutions t , so e will be a weakest link. Thus, we may assume that all of the candidate weakest links are flexible edges.

Now let s be the fixed solution for P_{opt} that assigns the earliest time to each temporal variable. According to lemma 1, there must also be another solution \hat{s} that differs from s on all of the candidate weakest links, and is as close as we like to s . In particular, we may choose \hat{s} so that the preference function is monotonic over the range of values between $s(e)$ and $\hat{s}(e)$ for every edge e . (Each convex function can be broken into monotonic segments.)

We may now divide the candidate weakest links e into three classes: (1) where $f(s(e)) > opt$, (2) where $f(s(e)) = opt < f(\hat{s}(e))$, and (3) where $f(s(e)) = opt = f(\hat{s}(e))$. For e in

²For simplicity of notation, we suppress the implied e subscript on f here and in what follows.




Figure 6: An STPP with no weakest link

class (1), note that, since f is continuous and \hat{s} can be made arbitrarily close to s , we may assume without loss of generality that $f(\hat{s}(e)) > \text{opt}$. Now suppose class (3) is empty. Then $f(\hat{s}(e)) > \text{opt}$ for each candidate weakest link e . Note that $f(\hat{s}(e)) > \text{opt}$ is also true for edges that are not candidate weakest links. Thus, $f(\hat{s}(e)) > \text{opt}$ for every edge e , which is a contradiction, because it means opt can be increased.

We conclude from this that the assumption that the third class is empty is untenable. Therefore, there is some edge e such that $f(s(e)) = f(\hat{s}(e)) = \text{opt}$. Recall that we constructed \hat{s} so that $s(e) \neq \hat{s}(e)$ and so that f is monotonic in the range between these two points. It follows that f is flat (constant) in this range. Since f is convex, it follows that it has its maximum in this range. (If a function is flat on part of its domain, and is higher either before or after that part, it cannot be convex.) Thus, the maximum of f is opt . We conclude that $f(t(e)) = \text{opt}$ for every WLO solution t , and so e is a weakest link. \square

The operations in the WLO+ algorithm preserve the convexity and continuity properties of the preference functions. Each stage of WLO+ repeats a WLO calculation. Thus, theorem 2 implies

Corollary 2.1 *Suppose P is an STPP with continuous domains and continuous convex preference functions. Then a weakest link is found at each iteration of WLO+.*

An example of why the existence result does not apply more generally to semi-convex functions is found in figure 6. The STPP in the figure contains two semi-convex step-like functions with optimal preference values associated with durations t and t' . Assume the STPP is minimal, and that the assignment $e = t, e' = t'$ is inconsistent. Then the highest possible chop point is p , and no weakest link exists, i.e., for neither e nor e' is it the case that, for every solution s , p is the value returned by the preference function associated with that constraint for the duration assigned by s .

Discussion. The proof of theorem 2 is in the classical setting of real-valued domains and continuous functions. In the more practical setting of a computer program where numbers are computed to a finite precision, some adjustments may need to be made. For example, lemma 1 depends on a counting argument that requires a modest amount of precision in the temporal representation.

Another issue involves the meaning of convexity in the case of a discrete approximation. Our approach is to treat a discretized function as corresponding to a piecewise linear function where the linear segments join successive points

on the discretized graph. We can then regard the discretized function as being convex if this piecewise linear function is convex. The results of this section are applicable to preference functions whose discretizations are convex in this sense. (Note that curves with overly gentle slopes, relative to the precision, may no longer be convex after they are discretized.)

An examination of the proof of theorem 2 shows that the weakest link exists under a somewhat less restrictive condition than convexity. It is enough (assuming semi-convexity) to require that plateaus can only occur at the global maximum of the preference function. This means, for example, that Gaussians can be allowed if their domains are trimmed to remove the flat parts of their (discretized) tails, provided that does not exclude all the solutions.

WLO+ has been implemented and tested on randomly generated problems, where each semi-convex preference function is a quadratic $ax^2 + bx + c$, with randomly selected parameters and $a \leq 0$. We compared the best solution found after applying WLO+ with the quality of the earliest solution found using the chop solver, using the utilitarian measure of quality (i.e., summing preference values). An average improvement of between 6 and 10% was observed, depending on constraint density (more improvement on lower density problems). Future research will focus on the application of WLO+ to the rover science planning domain.

6 Summary

This paper has presented a reformulation of problems in the optimization of temporal preferences using a generalization of Temporal CSPs. The practical context from which this effort arose is temporal decision-making in planning, where associated with domains representing temporal distances between events is a function expressing preferences for some temporal values over others. The work here extends previous work by overcoming limitations in the approach that arose when considerations of efficiency in reasoning with preferences resulted in coarseness in the evaluation procedure for global temporal assignments.

References

- [Bistarelli *et al.*, 1997] S. Bistarelli, U. Montanari, and F. Rossi. Semiring-based Constraint Solving and Optimization. *Journal of the ACM*, 44(2):201–236, March 1997.
- [Cormen *et al.*, 1990] T.H. Cormen, C.E. Leiserson, and R.L. Rivest. *Introduction to Algorithms*. MIT press, Cambridge, MA, 1990.
- [Dechter *et al.*, 1991] R. Dechter, I. Meiri, and J. Pearl. Temporal constraint networks. *Artificial Intelligence*, 49, 1991:61–95, 1991.
- [Khatib *et al.*, 2001] L. Khatib, P. Morris, R. Morris and F. Rossi. Temporal Reasoning about Preferences. In *Proceedings of IJCAI-01*. Morgan Kaufman, 1995.